

CONVEXITY OF A STOCHASTIC CONTROL FUNCTIONAL RELATED TO IMPORTANCE SAMPLING OF ITÔ DIFFUSIONS

HAN CHENG LIE^{†‡}

Abstract. We consider the problem of rare event importance sampling, where the random variable of interest is a path functional of an Itô diffusion computed up to the first exit from a d -dimensional bounded domain. Dupuis and Wang (*Ann. Appl. Probab.*, 15 (2005), pp. 1-38) studied the importance sampling problem by formulating it as a stochastic optimal control problem, where the value function is related to the conditional cumulant generating function of the random variable. In this paper, we show that the sufficient conditions for the value function to be twice-differentiable and α -uniformly Hölder continuous on the closure of the domain are also sufficient conditions for positive definiteness of the second variation of the control functional on the space of differentiable, α -uniformly Hölder continuous \mathbb{R}^d -valued feedback controls. We derive an expression for the second variation using Kazamaki's sufficient condition for L^q -boundedness of exponential martingales, and using Fredholm theory to prove the finiteness of the moment generating function of the first exit time over any bounded interval containing the origin. The strict convexity result suggests that one may be able to solve the corresponding Hamilton-Jacobi-Bellman boundary value problem in a dimension-robust way, by combining convex optimisation and Monte Carlo methods. We apply the result to analyse a gradient descent algorithm proposed by Hartmann and Schütte (*J. Stat. Mech. Theor. Exp.* (2012), P11004) for efficient rare event simulation.

Key words. Optimal stochastic control, Hamilton-Jacobi-Bellman equation, Feynman-Kac formula, elliptic boundary value problem, rare events, first exit time, importance sampling, convex optimisation, Monte Carlo methods

AMS subject classifications. 49M30, 60H30, 65C60, 82C80, 93E20

1. Introduction. The estimation of statistical properties of rare events is a common problem in computational physics. For example, the estimation of thermodynamic free energy differences between two equilibrium states is integral to the study of complex macromolecules in computational biophysics and drug design. Accordingly, many methods for rare event estimation have been proposed, especially in the case of free energy computations. Some methods, e.g. the adaptive biasing force method [5] or hyperdynamics [26, 27], involve driving the dynamical system of interest by applying an external force; in certain cases, this is achieved by changing the (potential) energy landscape on which the system lives, e.g. metadynamics [19]. Other methods, such as the Passerone-Parrinello method [22], study a variational problem and apply a least action principle, while others focus on improving the simulation of rare trajectories that dominate the statistical property of interest, e.g. the transition path sampling method [6] or the forward flux sampling method [1]. A mathematical treatment of free energy computations is given in [20].

The motivation of the present article is the analysis of a gradient descent method for solving a stochastic optimal control problem, proposed by Hartmann and Schütte [13]. This method shares some characteristics with the methods mentioned above: the method improves the collection of rare event statistics by tilting the potential in order to bias the path measure towards trajectories that dominate an exponential average (see, e.g. [15]), each modification to the potential corresponds to the application of a feedback control force, and a variational principle enters via the duality relationship

[†]Institut für Mathematik, Freie Universität Berlin, Berlin, Germany (hlie@math.fu-berlin.de).

[‡]This research was funded by a Ph.D. scholarship from the International Max-Planck Research School for Computational Biology and Scientific Computing of the Max Planck-Institut für molekulare Genetik, and by the Deutsche Forschungsgemeinschaft (DFG) through grant CRC 1114.

between free energy and relative entropy (see, e.g. [4]).

In this article, we describe sufficient conditions under which the control functional of the stochastic optimal control problem is strictly convex, by rigorously deriving an expression for the second variation of the control functional, and by showing that a deterministic lower bound on the path functional suffices to guarantee strict positive definiteness of the second variation. Strict convexity implies the well-posedness of the restriction of the stochastic optimal control problem to any finite-dimensional subspace of the set of admissible feedback control functions. In particular, in the limit of infinitely large sample size and infinitesimally small descent step size, it follows that the gradient descent method mentioned above converges to the unique minimiser corresponding to the best approximation in the finite-dimensional subspace. For the functionals and systems considered here, the strict convexity result renders the importance sampling problem amenable to other methods for solving convex optimisation problems that may be more efficient than gradient descent. This is important, because the connections between the optimisation problems considered here and certain elliptic Dirichlet boundary value problems imply that one can obtain numerical methods for solving Hamilton-Jacobi-Bellman boundary value problems that scale well with the spatial dimension.

Our initial reference for the idea of formulating an importance sampling problem as a stochastic optimal control problem is [13], which examined the problem of accelerating simulations of rare events by using the connection between risk-sensitive, exit time control and importance sampling [7] and the aforementioned free energy-relative entropy duality relationship [4]; the work [10] also studied stochastic optimal control formulations of importance sampling problems. The connections between importance sampling, large deviations, and differential games has been studied in [9], these ideas have been further developed for rare event simulation of multiscale diffusions [8, 25]. We shall not consider the large deviations aspect of rare event importance sampling. Instead, we shall view importance sampling as an optimisation problem defined over a set of measures; this idea has been applied to Markov decision processes in [3], and in the cross-entropy minimisation method introduced by Kullback [18]. Cross-entropy minimisations have been applied to importance sampling for rare event estimation by Rubenstein in the work [24], and recently by Asmussen, Kroese, and Rubenstein in [2]. In molecular dynamics, importance sampling and cross-entropy minimisation have been treated in [12], and Zhang et. al. [29].

The present work may differ from some of those mentioned above in the following ways. First, although we first consider the importance sampling problem, our primary concern is an analysis of the stochastic optimal control problem via the control functional. Second, our analysis is motivated by the desire to better understand what numerical methods may be applied to solve such problems efficiently in high-dimensional domains. Third, we employ functional analytic tools from the theory for classical solutions to elliptic partial differential equations, as opposed to large deviations techniques. Furthermore, since the approach we consider does not involve minimising the relative entropy alone, it differs from the cross-entropy minimisation methods mentioned earlier. To the best of our knowledge, no work to date has addressed the question of whether the control functional is convex over a suitably chosen function space. An advantage of our functional-analytic approach is that the sufficient conditions we specify are deterministic; such conditions may be easier to verify in practical situations.

2. Overview. Let $d \in \mathbb{N}$ be the dimension of the state space (configuration space) of a stochastic dynamical system whose paths belong to the canonical filtered probability space $(\Theta, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, P)$, where $\Theta = C([0, \infty); \mathbb{R}^d)$ and P is the standard Wiener measure. Consider an Itô diffusion, i.e. a stochastic dynamical system with continuous paths that satisfies the stochastic differential equation (SDE)

$$(2.1) \quad dX_t^u = (u - \nabla V)(X_t^u)dt + \sqrt{2\varepsilon}dB_t, \quad X_0^u = x,$$

where $u = (u_t)_{t \geq 0}$ denotes a predictable control process, $V \in C^1(\mathbb{R}^d; \mathbb{R})$ is an (potential) energy function that is bounded from below and that satisfies the growth conditions

$$(2.2a) \quad |V(x')| \leq K_0(1 + |x'|)$$

$$(2.2b) \quad |\nabla V(x')|^2 \leq K_1^2(1 + |x'|^2)$$

for some constants $K_0, K_1 > 0$ and for all $x' \in \mathbb{R}^d$, $\varepsilon > 0$ is a fixed temperature, and B is the standard Brownian motion with respect to P . In the context of computational physics and molecular dynamics, the SDE (2.1) is the overdamped Langevin equation that models the evolution of a large molecule in equilibrium with a heat bath composed of many smaller solvent molecules, where ∇V corresponds to the force field that arises from bonded and non-bonded interactions of the constituent atoms, and the white noise term models the effect of the heat bath on the molecule. The control u models the effect of a controllable external force that acts on the molecule, such as a pressure or electromagnetic field. Given u , the corresponding solution to (2.1) is a random variable $X^u : \Theta \rightarrow \Theta$. Letting $P^x(A) := P(\theta \in A | \theta_0 = x)$ denote the conditioning of P on the deterministic initial condition $x \in \mathbb{R}^d$, we then write the law of X^u , its conditioned version, and the conditional expectation of a path functional $\phi \in L^1(\Theta; \mathbb{R})$ as

$$(2.3) \quad \mu^u := P \circ (X^u)^{-1}, \quad \mu^{u,x} := P^x \circ (X^u)^{-1}, \quad E^{u,x}[\phi] = \int_{\Theta} \phi(\theta) \mu^{u,x}(d\theta).$$

We shall refer to $\mu^{0,x}$ as the ‘reference measure’, because all desired statistics shall defined with respect to $\mu^{0,x}$.

2.1. Importance sampling and stochastic optimal control. For a given path functional $W \in L^1(\Theta, \mu^{0,x}; \mathbb{R})$, we wish to estimate the statistics of W with respect to $\mu^{0,x}$. In the setting of importance sampling, an optimal importance sampling measure is any measure that yields a zero-variance estimator of the desired expected value. In our case, the optimal importance sampling measure solves the problem

$$(2.4) \quad \min \left\{ E^{u,x} \left[\left(\phi \frac{d\mu^0}{d\mu^u} - E^{u,x} \left[\phi \frac{d\mu^0}{d\mu^u} \right] \right)^2 \right] \mid \mu^u|_{\mathcal{F}_t} \ll \mu^0|_{\mathcal{F}_t}, \quad \mu^0|_{\mathcal{F}_t} \ll \mu^u|_{\mathcal{F}_t} \quad \forall t > 0 \right\}.$$

Thus, we seek to minimise the variance (with respect to $\mu^{u,x}$) of $\phi \frac{d\mu^0}{d\mu^u}$ over the set of measures μ^u that are mutually equivalent to μ^0 when restricted to \mathcal{F}_t for all $t > 0$. Recall that if the control process u' satisfies Kazamaki’s or Novikov’s condition, then the Radon-Nikodym derivative of $\mu^{u+u'}$ with respect to μ^u is a uniformly integrable exponential martingale with respect to $\mu^{u+u'}$, and the desired mutual equivalence

holds, by the Cameron-Martin-Girsanov theorem on change of measure. The Radon-Nikodym derivative of μ^u with respect to $\mu^{u+u'}$ is

$$(2.5) \quad \frac{d\mu^u}{d\mu^{u+u'}} \Big|_{\mathcal{F}_t} = \mathcal{E}(M^{-u'})_t := \exp \left[M_t^{-u'} - \frac{1}{2} \langle M^{-u'} \rangle_t \right]$$

where $M^{u'}$ is a continuous local martingale with respect to $(\mathcal{F}_t)_{t \geq 0}$ and $\langle M^{u'} \rangle$ is the associated quadratic variation process,

$$M_t^{u'} := \frac{1}{\sqrt{2\varepsilon}} \int_0^t u'_s dB_s, \quad \langle M^{u'} \rangle_t := \frac{1}{2\varepsilon} \int_0^t |u'_s|^2 ds.$$

By (2.5), we have the reweighting formula

$$(2.6) \quad E^{u,x} [W] = E^{u+u',x} \left[W \mathcal{E}(M^{-u'}) \right].$$

A useful mnemonic device for applying the reweighting formula (2.6) is that the sum of the control superscripts on the left-hand side equals the sum of the control superscripts on right-hand side: $u = (u + u') + (-u')$.

We now give some structure to the path functional of interest. Let $\Omega \subset \mathbb{R}^d$ be an open, bounded, and connected set, where every point on the boundary $\partial\Omega$ satisfies an exterior sphere condition. Define the first exit time of X^u from Ω by

$$(2.7) \quad \tau(X^u) := \inf \{t > 0 \mid X_t^u \notin \Omega\}.$$

Let $\kappa_r \in C(\Omega; [0, \infty))$ and $\kappa_t \in C(\partial\Omega; \mathbb{R})$. Define the path functional W by

$$(2.8) \quad W(X^u) := \int_0^{\tau(X^u)} \kappa_r(X_s^u) ds + \kappa_t(X_{\tau(X^u)}^u).$$

We shall refer to κ_r and κ_t as the ‘running cost’ and ‘terminal cost’ respectively. Let $\sigma > 0$ be fixed, and define the functions

$$(2.9a) \quad \psi(\sigma, x) := E^{0,x} [\exp(-\sigma W)]$$

$$(2.9b) \quad F(\sigma, x) := -\sigma^{-1} \log \psi(\sigma, x).$$

Observe that (2.9b) resembles Jarzynski’s equality [14] in statistical physics, which relates the thermodynamic free energy difference between two equilibrium states of a statistical-mechanical system to an exponential average of nonequilibrium work. In the statistical physics context, σ may be understood as an inverse temperature. Two important differences between the situation we consider and the situation considered by Jarzynski in [14] are that, in the former, the control force applied to the system is a function of the state and the control is applied up to a random stopping time, whereas in the latter, the control force is applied according to a preset schedule and over a deterministic time.

We wish to compute, for given σ and x , the value of $F(\sigma, x)$. By (2.9a), it suffices to compute $\psi(\sigma, x)$. Provided that u satisfies Novikov’s condition, we may combine (2.9a) and (2.6) in order to obtain

$$(2.10) \quad \psi(\sigma, x) = E^{u,x} [\exp(-\sigma W) \mathcal{E}(M^{-u})].$$

Thus, we may formulate the task of computing $\psi(\sigma, x)$ as an importance sampling problem of the form (2.4):

$$(2.11) \quad \min \left\{ E^{u,x} \left[\left(\exp(-\sigma W) \mathcal{E}(M^{-u}) - \psi(\sigma, x) \right)^2 \right] \mid \mathcal{E}(M^{-u}) \text{ uniformly integrable} \right\}.$$

In order to reformulate the importance sampling problem (2.11) as a stochastic optimal control problem, we define

$$(2.12) \quad K^{\sigma,u,\alpha} := W + (2\sigma)^{-1} \langle M^u \rangle_\tau + \alpha \sigma^{-1} M_\tau^u, \quad \alpha \in \{0, 1\}.$$

By (2.5) and (2.10), it follows that $\psi(\sigma, x) = E^{u,x} [\exp(-\sigma K^{\sigma,u,1})]$. For an arbitrary, adapted continuous local martingale M and stopping time T , we shall write

$$M_t^T := \begin{cases} M_{T \wedge t} & \{0 < T\} \\ 0 & \{0 = T\}. \end{cases}$$

Since the diffusion matrix is a strictly positive multiple of the identity matrix, it holds that $\tau \in L^1(\mu^{u,x})$ (see [16, Section 5.7, Lemma 7.4]). Thus, if $u \in L^\infty(\Omega)$, then it follows that $(M^u)^\tau$ is a $L^2(\mu^{u,x})$ -bounded martingale, since

$$\lim_{t \rightarrow \infty} E^{u,x} [((M^u)_t^\tau)^2] = E^{u,x} [\langle (M^u)^\tau \rangle_\infty] \leq \|u\|_{L^\infty}^2 E^{u,x}[\tau].$$

Since L^p -boundedness for $p > 1$ implies uniform integrability, we may apply the optimal stopping theorem to obtain $E^{u,x}[M_\tau^u] = 0$ and $E^{u,x}[K^{\sigma,u,1}] = E^{u,x}[K^{\sigma,u,0}]$. By the strict concavity of the logarithm, Jensen's inequality and (2.9b) yield

$$(2.13) \quad F(\sigma, x) \leq E^{u,x} [K^{\sigma,u,0}] =: \phi^{\sigma,x}(u)$$

where inequality holds if and only if the random variable $K^{\sigma,u,1}$ is $\mu^{u,x}$ -almost surely constant (see, e.g. [21, Theorem 5]). Thus, we have shown that we can compute $F(\sigma, x)$ by solving the stochastic optimal control problem

$$(2.14) \quad \min_{u \in \mathcal{U}} \phi^{\sigma,x}(u) \quad \text{subject to (2.1),}$$

where \mathcal{U} denotes the set of all previsible control processes. It was shown in [13] that there exists a unique, Markovian optimal control u_{opt}^σ given by

$$(2.15) \quad u_{\text{opt}}^\sigma(x') = -2\varepsilon\sigma \nabla_x F(\sigma, x') \quad \forall x' \in \Omega.$$

The link between the optimal control and importance sampling problems is evident from (2.15) since the optimal (zero-variance) importance sampling measure associated to u_{opt}^σ is defined in terms of F .

2.2. Connections to elliptic partial differential equations. In this section, we recall some fundamental results concerning the functions $\psi(\sigma, \cdot)$ and $F(\sigma, \cdot)$. We shall consider second order partial differential operators of the form

$$(2.16) \quad L := \sum_{i,j=1}^d a_{ij} \frac{\partial}{\partial x_i} \frac{\partial}{\partial x_j} + \sum_{i=1}^d b_i \frac{\partial}{\partial x_i} + c$$

where the functions $(a_{ij})_{i,j}, (b_i)_i, c \in C(\Omega; \mathbb{R})$. Following the nomenclature from stochastic analysis (see [16, Section 5.7A, Definition 7.1]), we shall say that L is ‘uniformly elliptic’ on Ω if there exists a constant $\lambda > 0$ such that

$$(2.17) \quad \lambda |\xi|^2 \leq \xi \cdot a \xi \quad \forall \xi \neq 0.$$

Recall that the linear Dirichlet boundary value problem is given by

$$(2.18a) \quad Lv(x) = f(x) \quad x \in \Omega,$$

$$(2.18b) \quad v(y) = \varphi(y) \quad y \in \partial\Omega,$$

for $f \in C(\Omega; \mathbb{R})$ and $\varphi \in C(\partial\Omega; \mathbb{R})$. Let u in (2.1) be an element of $C(\Omega; \mathbb{R}^d)$. Then the infinitesimal generator corresponding to the SDE (2.1) given by

$$(2.19) \quad \mathcal{A}^u := \varepsilon \Delta + (u - \nabla V) \cdot \nabla$$

is a second order, uniformly elliptic partial differential operator. Define the linear Dirichlet boundary value problem

$$(2.20a) \quad \mathcal{A}^0 \psi(\sigma, x) - \sigma \kappa_r(x) \psi(\sigma, x) = 0 \quad x \in \Omega$$

$$(2.20b) \quad \psi(\sigma, y) = \exp(-\sigma \kappa_t(y)) \quad y \in \partial\Omega.$$

We can rewrite (2.20) as a problem of the form (2.18), with $L := \mathcal{A}^0 - \sigma \kappa_r$ and $f = 0$ on Ω , and with $\varphi = \exp(-\sigma \kappa_t)$ on $\partial\Omega$. The following classical result establishes (2.9a) as the ‘Feynman-Kac formula’ for the solution to (2.20):

PROPOSITION 2.1. *If*

- (i) Ω is an open, bounded domain,
- (ii) the first-order coefficient $-\nabla V$ satisfies the growth conditions (2.2), and
- (iii) $\kappa_r \in C(\overline{\Omega}; [0, \infty))$ and $\kappa_t \in C(\partial\Omega; \mathbb{R})$,

then there exists a unique solution $\psi(\sigma, \cdot) \in C^2(\Omega; \mathbb{R}) \cap C(\overline{\Omega}; \mathbb{R})$ that solves (2.20), and $\psi(\sigma, \cdot)$ is of the form (2.9a).

Proof. See [16, Section 5.7, Proposition 7.2] and [16, Section 5.7, Lemma 7.4]. \square

Proposition 2.1 is important in establishing a link between the function $\psi(\sigma, \cdot)$ defined in (2.9a) and the linear Dirichlet boundary value problem (2.20). In §3, we shall require more regularity on the solution $\psi(\sigma, \cdot)$ in order to show L^2 -boundedness of exponential martingales. Since we will also work with continuous functions that are bounded on the compact set $\overline{\Omega}$, we also wish to know when the optimal control u_{opt}^σ defined in (2.15) is bounded on $\overline{\Omega}$. Recall that a bounded domain $\Omega \subset \mathbb{R}^d$ is of class $C^{k,\alpha}$ for $k \in \mathbb{N}_0$ and $0 \leq \alpha \leq 1$ if $\partial\Omega$ is the graph of a $C^{k,\alpha}$ function of $d-1$ coordinates [11, Section 6.2].

THEOREM 2.2. *Let $0 \leq \alpha \leq 1$ and L be uniformly elliptic. If*

- (i) Ω is an open, bounded, $C^{2,\alpha}$ domain,
- (ii) the coefficients of L and f in (2.18a) belong to $C^\alpha(\overline{\Omega})$,
- (iii) the zeroth-order coefficient of L is nonpositive, i.e. $c \leq 0$ on Ω , and
- (iv) $\kappa_t \in C^{2,\alpha}(\overline{\Omega})$,

then the Dirichlet problem (2.18) has a unique solution $v \in C^{2,\alpha}(\overline{\Omega})$.

Proof. See [11, Section 6.3, Theorem 6.14]. \square

By the definitions (2.9a) and (2.9b), it holds that $\psi(\sigma, \cdot) = \exp(-\sigma F(\sigma, \cdot))$. By the chain rule from ordinary calculus, we can formally transform the linear Dirichlet problem (2.20) into the nonlinear Dirichlet problem

$$(2.21a) \quad \mathcal{A}^0 F(\sigma, x) - \varepsilon \sigma |\nabla_x F(\sigma, x)|^2 + \kappa_r(x) = 0 \quad x \in \Omega$$

$$(2.21b) \quad F(\sigma, y) = \kappa_t(y) \quad y \in \partial\Omega,$$

where (2.21a) is the Hamilton-Jacobi-Bellman (HJB) equation of the stochastic optimal control problem (2.14). If there exists a unique solution $F(\sigma, \cdot)$ of (2.21), then it is the value function of the stochastic optimal control problem. In general, there are no classical solutions to (2.21), and one must search for a viscosity solution. Although viscosity solutions are interesting objects in their own right, our present interest pertains solely to classical solutions. For completeness, we shall consider sufficient conditions for which there exists a classical solution to the nonlinear Dirichlet problem (2.21).

Let $\Gamma := \Omega \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^{d \times d}$, $\gamma = (x, z, p, r) \in \Gamma$ be arbitrary, and $\mathcal{H} : \Gamma \rightarrow \mathbb{R}$ be defined by

$$(2.22) \quad \mathcal{H}(x, z, p, r) := \varepsilon \left(\sum_i r_{ii} \right) - \nabla V(x) \cdot p - \varepsilon \sigma |p|^2 + \kappa_r(x).$$

From (2.22), it follows that \mathcal{H} is constant with respect to z , concave with respect to p , and linear with respect to r . Note that we can rewrite (2.21a) as

$$\mathcal{H}(x, F(\sigma, x), \nabla_x F(\sigma, x), \nabla_x^2 F(\sigma, x)) = 0 \quad x \in \Omega.$$

We have

THEOREM 2.3. *Let \mathcal{H} be defined as in (2.22).*

(i) *If Ω is a bounded domain satisfying an exterior sphere condition at each boundary point, and*

(ii) *$\nabla V \in C^2(\Omega; \mathbb{R}^d)$ and $\kappa_r \in C^2(\Omega; \mathbb{R})$,*

then the classical Dirichlet problem (2.21) is uniquely solvable in $C^2(\Omega; \mathbb{R}) \cap C(\bar{\Omega}; \mathbb{R})$ for any terminal cost $\kappa_t \in C(\partial\Omega; \mathbb{R})$.

Proof. The proof follows from an application of [11, Section 17.5, Theorem 17.17]. Therefore, we must verify that the assumptions of that result hold for the operator \mathcal{H} defined in (2.22). The assumptions of the present theorem guarantee that $\mathcal{H} \in C^2(\Gamma)$, and that \mathcal{H} is constant with respect to z , concave with respect to p , and linear with respect to r . Let $\mathcal{H}_{ij}(x, z, p, r) := \frac{\partial \mathcal{H}}{\partial r_{ij}}(x, z, p, r)$. It remains to verify that the structure conditions (see [11, Section 17.5, Eq. (17.53)])

$$(2.23a) \quad 0 < \lambda |\xi|^2 \leq \sum_{i,j} \mathcal{H}_{ij} \xi_i \xi_j \leq \Lambda |\xi|^2$$

$$(2.23b) \quad |\mathcal{H}_p|, |\mathcal{H}_z|, |\mathcal{H}_{rx}|, |\mathcal{H}_{px}|, |\mathcal{H}_{zx}| \leq \mu \lambda,$$

$$(2.23c) \quad |\mathcal{H}_x|, |\mathcal{H}_{xx}| \leq \mu \lambda (1 + |p| + |r|),$$

are satisfied for nonzero $\xi \in \mathbb{R}^d$, arbitrary $(x, z, p, r) \in \Gamma$, a function λ that is non-increasing in $|z|$, and functions Λ and μ that are nondecreasing in $|z|$. By (2.22), condition (2.23a) holds with $\lambda = \Lambda = \varepsilon$ since $\mathcal{H}_{ij} = \delta_{ij} \varepsilon$. Furthermore, \mathcal{H}_z and \mathcal{H}_{rx} vanish everywhere, and

$$(2.24a) \quad \frac{\partial \mathcal{H}}{\partial p_i}(x, z, p, r) = -\frac{\partial V}{\partial x_i}(x) - 2\varepsilon \sigma p_i,$$

$$(2.24b) \quad \frac{\partial^2 \mathcal{H}}{\partial p_i \partial x_j}(x, z, p, r) = -\frac{\partial^2 V}{\partial x_i \partial x_j}(x),$$

$$(2.24c) \quad \frac{\partial \mathcal{H}}{\partial x_i}(x, z, p, r) = \frac{\partial}{\partial x_i} (\kappa_r(x) - \nabla V(x) \cdot p),$$

$$(2.24d) \quad \frac{\partial^2 \mathcal{H}}{\partial x_i \partial x_j}(x, z, p, r) = \frac{\partial^2}{\partial x_i \partial x_j} (\kappa_r(x) - \nabla V(x) \cdot p).$$

Since neither of the right-hand sides of (2.24) depends on $|z|$, we may define $\mu : \Gamma \rightarrow \mathbb{R}$ to be any function that satisfies (2.23b)–(2.23c). \square

Note that we may generalise Theorem 2.3 by weakening the requirement (ii), i.e. by only requiring that ∇V and κ_r be twice weakly differentiable functions of the spatial variable x and interpreting the structure conditions (2.23) in the weak sense [11, Chapter 17, Problem 17.4].

Theorem 2.2 may be generalised to larger classes of domains that need not be bounded (see [11, Section 6.3, Theorem 6.13] and the notes at the end of [11, Chapter 6]). Note that even if the zeroth-order coefficient of the operator L fails to satisfy the nonpositivity condition, solutions of (2.18) still belong to $C^{k+2,\alpha}(\overline{\Omega})$ for $k \in \mathbb{N}_0$, provided that the domain is $C^{k+2,\alpha}$, the boundary function belongs to $C^{k+2,\alpha}(\overline{\Omega})$, and the coefficients of the uniformly elliptic operator L belong to $C^{k,\alpha}(\overline{\Omega})$. However, uniqueness of solutions no longer holds [11, Section 6.4, Theorem 6.19].

3. Variational analysis of the control functional. In this section, we derive expressions for the first variation and the second variation of the control functional with respect to admissible perturbations of the feedback control function. In Lemma 3.2, we adapt an argument due to Gilbarg and Trudinger, and use Fredholm theory in order to assert that the moment generating function of the first exit time defined in (2.7) is finite over any open, bounded interval containing the origin. We then use Lemma 3.2 to show that Kazamaki's sufficient condition (3.3) for L^q -boundedness of exponential martingales holds for the case that $q > 1$. Together with the Cauchy-Schwarz inequality and Lebesgue's dominated convergence theorem, L^q -boundedness yields the expressions for the first and second variations. The main result of this section, Theorem 3.14, shows the strict convexity of the control functional by proving that the second variation is a positive definite functional on the Banach space $C^{1,\alpha}(\overline{\Omega}; \mathbb{R}^d)$ whenever the work path functional W defined in (2.8) is bounded from below. Recall that $\sigma > 0$ is a fixed parameter. All the results in this section, unless otherwise stated, shall be based upon the following

ASSUMPTION 3.1. *Let $0 \leq \alpha \leq 1$. It holds that*

- (i) *the domain Ω is bounded and of class $C^{2,\alpha}$,*
- (ii) *the gradient of the potential $\nabla V \in C^\alpha(\overline{\Omega}; \mathbb{R}^d) \cap W^2(\Omega; \mathbb{R}^d)$,*
- (iii) *the running cost $\kappa_r \in C^\alpha(\overline{\Omega}; [0, \infty)) \cap W^2(\Omega; \mathbb{R})$, and*
- (iv) *the terminal cost $\kappa_t \in C^{2,\alpha}(\overline{\Omega}; \mathbb{R})$.*

By Theorem 2.2, there exists a unique solution to (2.18) in $C^{2,\alpha}(\overline{\Omega}; \mathbb{R})$ with $f = 0$ on Ω and $\varphi = \exp(-\sigma\kappa_t)$ on $\partial\Omega$. Proposition 2.1 guarantees that the unique solution is described by the Feynman-Kac formula (2.9a). It follows from (2.9a) and Jensen's inequality that $\exp(E^{0,x}[-\sigma W]) \leq \psi(\sigma, x)$; since finiteness of $\psi(\sigma, x)$ on $\overline{\Omega}$ implies finiteness of $E^{0,x}[-\sigma W]$, it follows that $\psi(\sigma, x)$ is bounded away from zero on $\overline{\Omega}$. On the other hand, the maximum principle implies that $\psi(\sigma, x)$ is bounded from above. Thus, given the relation (2.9b), it follows that $F(\sigma, \cdot) \in C^{2,\alpha}(\overline{\Omega}; \mathbb{R})$. By Theorem 2.3, $F(\sigma, \cdot)$ is a classical solution of the HJB problem (2.21), and the optimal control belongs to $C^{1,\alpha}(\overline{\Omega}; \mathbb{R}^d)$.

3.1. L^q -bounded exponential martingales. The first result below is of crucial importance:

LEMMA 3.2. *Suppose that Assumption 3.1 holds. If $u \in C^\alpha(\overline{\Omega})$, then for all $r > 0$, there exists $\sigma_p > r$ such that $E^{u,x}[\exp(\sigma_p \tau)]$ is finite for all $x \in \overline{\Omega}$.*

Proof. We prove the statement using Fredholm theory, following the argument outlined at the end of [11, Section 6.3]. Observe that, for an arbitrary $\sigma_n < 0$, $E^{u,x}[\exp(\sigma_n \tau)]$ uniquely solves the Feynman-Kac boundary value problem (2.20) with

the choice of $-\sigma = \sigma_n < 0$, $\kappa_r = 1$ and $\kappa_t = 0$. Since κ_r and κ_t are constant, they may be extended to all of $\bar{\Omega}$. Therefore, we may conclude from Theorem 2.2 that $E^{u,x}[\exp(\sigma_n \tau)]$ belongs to $C^{2,\alpha}(\bar{\Omega}; \mathbb{R})$. We rewrite this Feynman-Kac boundary value problem as an elliptic boundary value problem of the form (2.18):

$$(3.1a) \quad L^{u,\sigma_n} v(x) := \mathcal{A}^u v(x) + \sigma_n v(x) = 0, \quad x \in \Omega,$$

$$(3.1b) \quad v(y) = 1, \quad y \in \partial\Omega.$$

Since all the coefficients of L^{u,σ_n} belong to $C^\alpha(\bar{\Omega})$, the operator L defines a closed map from coefficients that belong to $C^\alpha(\bar{\Omega})$ to solutions that belong to $C^{2,\alpha}(\bar{\Omega})$. By [11, Theorem 5.5], it follows that there is a countable, discrete set $\Sigma \subset \mathbb{R}$, such that if $\sigma \in \Sigma$, then the homogeneous problem

$$(3.2a) \quad L^{u,\sigma_n-\sigma} v(x) = L^{u,\sigma_n} v(x) - \sigma v(x) = 0, \quad x \in \Omega,$$

$$(3.2b) \quad v(y) = 0, \quad y \in \partial\Omega$$

has nontrivial solutions. Therefore, for $\sigma \notin \Sigma$, the homogeneous problem (3.2) has only the trivial solution $v \equiv 0$. By the Fredholm alternative (see [11, Section 6.3, Theorem 6.15]), it follows that the corresponding inhomogeneous problem obtained by replacing (3.2b) with (3.1b) has a unique solution in $C^{2,\alpha}(\bar{\Omega})$. By the extreme value theorem and compactness of $\bar{\Omega}$, the unique solution is finite on $\bar{\Omega}$. By Itô's formula, the unique solution admits the representation $v(x) = E^{u,x}[\exp((\sigma_n - \sigma)\tau)]$. Now let $r > 0$ be arbitrary. Since Σ is discrete, there exists a $\sigma \in (-\infty, 0) \cap \Sigma^c$ such that $\sigma_p := \sigma_n - \sigma > r$, and the conclusion follows. \square

COROLLARY 3.3. *For all $x \in \bar{\Omega}$, $\sigma_n < 0$ and $r > 0$, there exists a closed interval $[\sigma_n, \sigma_p]$ with $\sigma_p > r$ such that the moment generating function defined by $E^{u,x}[\exp(\sigma\tau)]$ is finite for all $\sigma \in [\sigma_n, \sigma_p]$.*

Proof. Let $x \in \bar{\Omega}$ be arbitrary. For any $\sigma_n < 0$, $E^{u,x}[\exp(\sigma_n \tau)] \leq 1$ by monotonicity of the integral. By Lemma 3.2, there exists a $\sigma_p > r$ such that $E^{u,x}[\exp(\sigma_p \tau)]$ is finite. For arbitrary $\sigma_n < \sigma < \sigma_p$, there exists a $0 < \theta < 1$ such that $\sigma = \theta\sigma_n + (1-\theta)\sigma_p$. By the strict convexity of the exponential, we have

$$\exp(\sigma\tau) < \theta \exp(\sigma_n \tau) + (1-\theta) \exp(\sigma_p \tau).$$

Taking expectations with respect to $\mu^{u,x}$ yields the desired conclusion. \square

THEOREM 3.4. *Let τ be an arbitrary stopping time, and let $\sigma > 0$ be such that $E[\exp(\sigma\tau)]$ is finite. If for some $1 < p < \infty$, it holds that $\varphi \in L^\infty(\Omega)$ satisfies*

$$(3.3) \quad \|\varphi\|_{L^\infty(\Omega)}^2 \leq 2\sigma \frac{(\sqrt{p}-1)^2}{p},$$

then the stopped exponential martingale $\mathcal{E}((M^\varphi)^\tau)$ is an L^q -bounded martingale for $q^{-1} = 1 - p^{-1}$ and $E[\mathcal{E}(M^\varphi)^q]$ is finite.

Proof. Let p and q be a pair of Hölder conjugate exponents. Recall [17, Section 1.2, Theorem 1.5]: for a given continuous local martingale M , if

$$(3.4) \quad \sup \left\{ E \left[\exp \left(\frac{1}{2} \frac{\sqrt{p}}{\sqrt{p}-1} M_T \right) \right] \mid T \text{ a bounded stopping time} \right\} < \infty,$$

then the exponential martingale $\mathcal{E}(M)$ is L^q -bounded. Let T be an arbitrary bounded stopping time. Since it follows from (2.5) that

$$\exp \left(\frac{1}{2} M_T \right) = (\mathcal{E}(M)_T)^{1/2} \exp \left(\frac{1}{4} \langle M \rangle_T \right),$$

it holds by the Cauchy-Schwarz inequality that

$$(3.5) \quad E \left[\exp \left(\frac{1}{2} M_T \right) \right] \leq E \left[\exp \left(\frac{1}{2} \langle M \rangle_T \right) \right]^{1/2}.$$

Combining the inequality (3.5) with (3.4), we therefore obtain that

$$(3.6) \quad \sup \left\{ E \left[\exp \left(\frac{p}{2(\sqrt{p}-1)^2} \langle M \rangle_T \right) \right] \mid T \text{ a bounded stopping time} \right\} < \infty$$

is a sufficient condition for the L^q -boundedness of $\mathcal{E}(M)$. By the assumption that $\varphi \in L^\infty(\Omega)$, it holds that

$$\langle (M^\varphi)^\tau \rangle_T \leq \|\varphi\|_{L^\infty(\Omega)}^2 (\tau \wedge T).$$

Since (3.3) is equivalent to

$$\frac{p}{2(\sqrt{p}-1)^2} \|\varphi\|_{L^\infty(\Omega)}^2 \leq \sigma,$$

and since $E[\exp(\sigma\tau)]$ is finite, it follows that (3.6) is satisfied, by the monotonicity of the exponential and of the integral. Therefore $\mathcal{E}(M^\varphi)^\tau$ is L^q -bounded. By the definition of L^q -boundedness, it follows that $\mathcal{E}(M^\varphi)_\tau \in L^q(\mu^{u,x})$, as desired. \square

The next result combines Lemma 3.2 and Theorem 3.4 to show that, for any $\varphi \in L^\infty(\Omega; \mathbb{R}^d)$, it holds that $\mathcal{E}((M^\varphi)^\tau)$ is $L^q(\mu^{u,x})$ -bounded for any $q > 1$.

COROLLARY 3.5. *Let τ be defined as in (2.7) and suppose that Assumption 3.1 holds. If $u \in C^\alpha(\overline{\Omega})$ and $\varphi \in L^\infty(\Omega)$, then*

- (i) $\mathcal{E}(M^\varphi)^\tau$ is $L^q(\mu^{u,x})$ -bounded for all $q > 1$, and
- (ii) for any $1 \leq r < \infty$, $E^{u,x}[(M^\varphi)_\tau^r]$ is finite, and for any $0 < s < \infty$, $E^{u,x}[\sup_{0 \leq t \leq \tau} |M_t^\varphi|^s]$ is finite.

Proof. We first prove statement (i). Let p be the Hölder conjugate of q , and define

$$r_p := \frac{p}{2(\sqrt{p}-1)^2} \|\varphi\|_{L^\infty(\Omega)}^2.$$

Since $u \in C^\alpha(\overline{\Omega})$, there exists a $\sigma_p > r_p$ so that $E^{u,x}[\exp(\sigma_p \tau)]$ is finite for all $x \in \overline{\Omega}$, by Lemma 3.2. Hence, φ satisfies (3.3), and the conclusion follows by Theorem 3.4.

To prove statement (ii), observe that (3.6) implies that $E[(M^\varphi)_\tau^r]$ is finite for all $r \in \mathbb{N}$, and hence for all $1 \leq r < \infty$, by Jensen's inequality. The second assertion of statement (ii) follows from this observation and the Burkholder-Davis-Gundy inequalities [23, Chapter IV, §4, Theorem 4.1]. \square

The next result shows that Assumption 3.1 suffices to guarantee that, for bounded controls u , the quantity $E^{u,x}[\exp(-\sigma W)]$ is finite for all $x \in \overline{\Omega}$.

LEMMA 3.6. *Let W be defined as in (2.8) and suppose that Assumption 3.1 holds. If $u \in L^\infty(\Omega)$, then $E^{u,x}[\exp(-\sigma W)]$ is finite, and $W \in L^n(\mu^{u,x})$ for all $n \in \mathbb{N}$.*

Proof. Since the constant 0 belongs to $C^\alpha(\overline{\Omega})$, and since we assumed that $u \in L^\infty(\Omega)$, we may apply Corollary 3.5 to assert that $\mathcal{E}(M^u)^\tau$ is $L^2(\mu^{0,x})$ -bounded. By the reweighting formula (2.6) and the Cauchy-Schwarz inequality,

$$E^{u,x}[\exp(-\sigma W)] \leq (E^{0,x}[\exp(-2\sigma W)] E^{0,x}[\mathcal{E}(M^u)_\tau^2])^{1/2},$$

where Theorem 2.2 implies that the first term on the right-hand side is finite and $L^2(\mu^{0,x})$ -boundedness of $\mathcal{E}(M^u)_\tau$ implies that the second term on the right-hand side

is finite. This proves the first conclusion. The second conclusion follows from the series expansion of the exponential. \square

The next result of this section will be useful later and follows from statement (ii) of Corollary 3.5 and Lemma 3.6.

COROLLARY 3.7. *Suppose that Assumption 3.1 holds. If $\sigma > 0$ and $u \in C^\alpha(\overline{\Omega})$, then $K^{\sigma,u,0}, K^{\sigma,u,1} \in L^n(\mu^{u,x})$ for all $n \in \mathbb{N}$.*

3.2. Convergence lemma. By the definition (2.5) of the exponential martingale, it holds that for $0 \neq \delta \in \mathbb{R}$, a Borel-measurable function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^d$, and a stopping time τ ,

$$\mathcal{E}(M^{\delta\varphi})_\tau = \exp \left[M_\tau^{\delta\varphi} - \frac{1}{2} \langle M^{\delta\varphi} \rangle_\tau \right] = \exp \left[\delta \left(M_\tau^\varphi - \frac{\delta}{2} \langle M^\varphi \rangle_\tau \right) \right].$$

The series expansion of the exponential thus yields

$$(3.7) \quad \mathcal{E}(M^{\delta\varphi})_\tau = 1 + \sum_{m=1}^{\infty} \frac{\delta^m (M_\tau^\varphi - \delta 2^{-1} \langle M^\varphi \rangle_\tau)^m}{m!}$$

and hence

$$(3.8) \quad \frac{\mathcal{E}(M^{\delta\varphi})_\tau - 1}{\delta} - M_\tau^\varphi = -\frac{\delta}{2} \langle M^\varphi \rangle_\tau + \sum_{m=2}^{\infty} \frac{\delta^{m-1} (M_\tau^\varphi - \delta 2^{-1} \langle M^\varphi \rangle_\tau)^m}{m!}.$$

We have the following convergence result:

LEMMA 3.8. *Suppose that Assumption 3.1 holds, that $u \in C^\alpha(\overline{\Omega})$, and $\varphi \in L^\infty(\Omega)$. Then for any nonzero sequence $(\delta_n)_n$ satisfying $\delta_n \rightarrow 0$ as $n \rightarrow \infty$, it holds that $\mathcal{E}(M^{\delta_n\varphi})_\tau \in L^2(\mu^{u,x})$ for all $n \in \mathbb{N}$, and*

$$(3.9) \quad \lim_{n \rightarrow \infty} E^{u,x} \left[(\mathcal{E}(M^{\delta_n\varphi})_\tau - 1)^2 \right] = 0,$$

$$(3.10) \quad \lim_{n \rightarrow \infty} E^{u,x} \left[\left(\frac{\mathcal{E}(M^{\delta_n\varphi})_\tau - 1}{\delta_n} - M_\tau^\varphi \right)^2 \right] = 0.$$

Proof. Without loss of generality, we may assume that the sequence $(|\delta_n|)_n$ is bounded by $0 < r < 1$, so that

$$(3.11) \quad \|\delta_n \varphi\|_{L^\infty(\Omega)} \leq r \|\varphi\|_{L^\infty(\Omega)} < \|\varphi\|_{L^\infty(\Omega)}.$$

Therefore, for all $n \in \mathbb{N}$, $\delta_n \varphi \in L^\infty(\Omega)$, and by Corollary 3.5, it follows that $\mathcal{E}(M^{\delta_n\varphi})_\tau \in L^2(\mu^{u,x})$ for all $n \in \mathbb{N}$. By Hölder's inequality, we have (see the proof of [17, Section 1.2, Theorem 1.5])

$$E^{u,x} [\mathcal{E}(M^{\delta_n\varphi})_\tau^2] \leq E^{u,x} \left[\exp \left(\frac{\sqrt{2}}{2(\sqrt{2}-1)} M_\tau^{\delta_n\varphi} \right) \right]^{2/(\sqrt{2}+1)}.$$

Applying (3.5) and (3.11) to the inequality above yields

$$(3.12) \quad \begin{aligned} E^{u,x} [\mathcal{E}(M^{\delta_n\varphi})_\tau^2] &\leq E^{u,x} \left[\exp \left(\frac{1}{(\sqrt{2}-1)^2} \|\delta_n \varphi\|_{L^\infty(\Omega)}^2 \tau \right) \right]^{1/(\sqrt{2}+1)} \\ &< E^{u,x} \left[\exp \left(\frac{1}{(\sqrt{2}-1)^2} \|\varphi\|_{L^\infty(\Omega)}^2 \tau \right) \right]^{1/(\sqrt{2}+1)}. \end{aligned}$$

By Corollary 3.3, it follows that the collection of random variables $(\mathcal{E}(M^{\delta_n \varphi})_\tau)_n$ is bounded in $L^2(\mu^{u,x})$. We can therefore apply Lebesgue's dominated convergence theorem and the identity (3.7) in order to obtain

$$\lim_{n \rightarrow \infty} E^{u,x} \left[(\mathcal{E}(M^{\delta_n \varphi})_\tau - 1)^2 \right] = E^{u,x} \left[\lim_{n \rightarrow \infty} (\mathcal{E}(M^{\delta_n \varphi})_\tau - 1)^2 \right] = 0,$$

which proves (3.9). To prove (3.10), we first observe that by (3.7), the random series

$$(3.13) \quad R(\delta_n) := (\mathcal{E}(M^{\delta_n \varphi})_\tau - 1)^2 - \delta_n^2 (M_\tau^\varphi)^2 \in L^1(\mu^{u,x})$$

consists of terms that are cubic or higher in δ_n , so that $\lim_{n \rightarrow \infty} \delta_n^{-2} E^{u,x}[R(\delta_n)] = 0$. Then there exists some constant $C(r) > 0$ so that

$$(3.14) \quad \begin{aligned} E^{u,x} \left[\left(\frac{\mathcal{E}(M^{\delta_n \varphi})_\tau - 1}{\delta_n} \right)^2 \right] &= E^{u,x} [\langle M^\varphi \rangle_\tau] + C(r) \\ &\leq E^{u,x} [\|\varphi\|_{L^\infty(\Omega)}^2 \tau] + C(r). \end{aligned}$$

Thus the sequence $(\delta_n^{-1}(\mathcal{E}(M^{\delta_n \varphi})_\tau - 1))_n$ is bounded in $L^2(\mu^{u,x})$. Since M_τ^φ belongs to $L^2(\mu^{u,x})$ and since $L^2(\mu^{u,x})$ is a vector space, it follows that the sequence $(\delta_n^{-1}(\mathcal{E}(M^{\delta_n \varphi})_\tau - 1) - M_\tau^\varphi)_n$ is bounded in $L^2(\mu^{u,x})$. We can then apply Lebesgue's dominated convergence theorem and the identity (3.8) to obtain

$$\lim_{n \rightarrow \infty} E^{u,x} \left[\left(\frac{\mathcal{E}(M^{\delta_n \varphi})_\tau - 1}{\delta_n} - M_\tau^\varphi \right)^2 \right] = E^{u,x} \left[\lim_{n \rightarrow \infty} \left(\frac{\mathcal{E}(M^{\delta_n \varphi})_\tau - 1}{\delta_n} - M_\tau^\varphi \right)^2 \right] = 0,$$

as desired. \square

3.3. First variation of control functional. In order to derive an expression for the first variation of the control functional ϕ defined in (2.13), we shall proceed by showing the convergence in $L^1(\mu^{u,x})$ of a sequence of finite differences of $\phi^{\sigma,x}$. For a given sequence $(\delta_n)_n$ and function φ that satisfy the conditions of Lemma 3.8, we define the finite difference random variables by

$$(3.15) \quad h_n^{\sigma,u,\varphi} := K^{\sigma,u,0} \frac{\mathcal{E}(M^{\delta_n \varphi})_\tau - 1}{\delta_n} + \frac{1}{\sigma} \left(\langle M^u, M^\varphi \rangle_\tau + \frac{\delta_n}{2} \langle M^\varphi \rangle_\tau \right) \mathcal{E}(M^{\delta_n \varphi})_\tau.$$

We show that the $(h_n^{\sigma,u,\varphi})_n$ are indeed the appropriate finite differences in the next

LEMMA 3.9. *Suppose that Assumption 3.1 holds, that $u \in C^\alpha(\overline{\Omega})$, $\varphi \in L^\infty(\Omega)$ and $(\delta_n)_n$ be an arbitrary nonzero sequence satisfying $\delta_n \rightarrow 0$ as $n \rightarrow \infty$. Let $(h_n^{\sigma,u,\varphi})_n$ be the corresponding sequence given by (3.15). Then $(h_n^{\sigma,u,\varphi})_n \subset L^1(\mu^{u,x})$ and*

$$\frac{\phi^{\sigma,x}(u + \delta_n \varphi) - \phi^{\sigma,x}(u)}{\delta_n} = E^{u,x} [h_n^{\sigma,u,\varphi}] \quad \forall n \in \mathbb{N}.$$

Proof. In order to show that $\{h_n^{\sigma,u,\varphi} | n \in \mathbb{N}\} \subset L^1(\mu^{u,x})$, we shall apply the Cauchy-Schwarz inequality individually to the three summands on the right-hand side of (3.15). Recall the inequality of arithmetic and geometric means,

$$(3.16) \quad |ab| \leq 2^{-1}(a^2 + b^2),$$

with equality holding if and only if $a = b$, and the Kunita-Watanabe inequality

$$(3.17) \quad \langle M^\varphi, M^\zeta \rangle_\tau \leq (\langle M^\varphi \rangle_\tau \langle M^\zeta \rangle_\tau)^{1/2}$$

that holds almost surely, with equality holding almost surely if and only if $\varphi = \zeta$ except on a set of Lebesgue measure zero. The stated assumptions guarantee that $W \in L^2(\mu^{u,x})$, by Lemma 3.6; that the sequences $(\delta_n^{-1}(\mathcal{E}(M^{\delta_n \varphi})_\tau - 1))_n$ and $(\mathcal{E}(M^{\delta_n \varphi})_\tau)_n$ are bounded sequences in $L^2(\mu^{u,x})$, by (3.14) and (3.12); and that integer moments of quadratic variation processes belong to $L^2(\mu^{u,x})$, by Corollary 3.5 (ii). Since $K^{\sigma,u,0} \in L^2(\mu^{u,x})$ by Corollary 3.7 and

$$(3.18) \quad E^{u,x} [|\langle M^u, M^\varphi \rangle_\tau|^2] \leq E^{u,x} [\langle M^u \rangle_\tau + \langle M^\varphi \rangle_\tau^2]$$

by (3.16) and (3.17), it follows that $\{h_n | n \in \mathbb{N}\} \subset L^1(\mu^{u,x})$, thus proving the first statement.

To prove the second statement, observe that by Lemma 3.8, it follows that each $\mathcal{E}(M^{\delta_n \varphi})$ is a $L^2(\mu^{u,x})$ -bounded exponential martingale. Since L^q -boundedness for $q > 1$ implies uniform integrability, we can apply the reweighting formula (2.6) to express $\phi^{\sigma,x}(u + \delta_n \varphi)$ as an expectation with respect to $\mu^{u,x}$. By the definition of the quadratic variation, we have for any stopping time τ that

$$(3.19) \quad \langle M^{u+\delta_n \varphi} \rangle_\tau = \langle M^u \rangle_\tau + \delta_n^2 \langle M^\varphi \rangle_\tau + 2\delta_n \langle M^u, M^\varphi \rangle_\tau.$$

By the reweighting formula and (3.19) we therefore have

$$\phi^{\sigma,x}(u + \delta_n \varphi) = E^{u,x} \left[\left(W + \frac{\langle M^u \rangle_\tau}{2\sigma} + \frac{\delta_n^2}{2\sigma} \langle M^\varphi \rangle_\tau + \frac{\delta_n}{\sigma} \langle M^u, M^\varphi \rangle_\tau \right) \mathcal{E}(M^{\delta_n \varphi})_\tau \right].$$

Taking differences yields the desired conclusion. \square

We now define the limiting random variable to be

$$(3.20) \quad h^{\sigma,u,\varphi} := K^{\sigma,u,0} M_\tau^\varphi + \sigma^{-1} \langle M^u, M^\varphi \rangle_\tau,$$

and show that the expected value of $h^{\sigma,u,\varphi}$ is equal to the first variation of $\phi^{\sigma,x}(u)$ in the direction of φ :

THEOREM 3.10. *If the hypotheses of Lemma 3.9 are satisfied, then $h \in L^1(\mu^{u,x})$, h_n converges to h in $L^1(\mu^{u,x})$, and the first variation of $\phi^{\sigma,x}(u)$ along φ satisfies*

$$(3.21) \quad \Phi^{\sigma,x}(u; \varphi) := \lim_{n \rightarrow \infty} \frac{\phi^{\sigma,x}(u + \delta_n \varphi) - \phi^{\sigma,x}(u)}{\delta_n} = E^{u,x} [h^{\sigma,u,\varphi}].$$

Proof. The third conclusion follows from the second, since

$$(3.22) \quad |E^{u,x} [h^{\sigma,u,\varphi} - h_n^{\sigma,u,\varphi}]| \leq E^{u,x} [|h^{\sigma,u,\varphi} - h_n^{\sigma,u,\varphi}|].$$

The first conclusion, that $h^{\sigma,u,\varphi} \in L^1(\mu^{u,x})$, follows from the observation that

$$E^{u,x} [|h^{\sigma,u,\varphi}|] \leq E^{u,x} [|K^{\sigma,u,0} M_\tau^\varphi|] + \sigma^{-1} E^{u,x} [|\langle M^u, M^\varphi \rangle_\tau|],$$

and by using the Cauchy-Schwarz inequality, Corollary 3.7, and (3.18).

It remains to show the convergence in $L^1(\mu^{u,x})$ of $h_n^{\sigma,u,\varphi}$ to $h^{\sigma,u,\varphi}$. We apply the triangle inequality to obtain

$$(3.23) \quad |h_n^{\sigma,u,\varphi} - h^{\sigma,u,\varphi}| \leq \left| K^{\sigma,u,0} \left(\frac{\mathcal{E}(M^{\delta_n \varphi})_\tau - 1}{\delta_n} - M_\tau^\varphi \right) \right| + \frac{1}{\sigma} (|\langle M^u, M^\varphi \rangle_\tau (\mathcal{E}(M^{\delta_n \varphi})_\tau - 1)| + \frac{|\delta_n|}{2\sigma} \langle M^\varphi \rangle_\tau \mathcal{E}(M^{\delta_n \varphi})_\tau).$$

Recall that $(\mathcal{E}(M^{\delta_n \varphi})_\tau)_n$ is a bounded sequence in $L^2(\mu^{u,x})$, by (3.12). Therefore, by the Cauchy-Schwarz inequality and Lebesgue's dominated convergence theorem, the third term on the right-hand side converges to zero in $L^1(\mu^{u,x})$. By the Cauchy-Schwarz inequality, the bound (3.18), and the convergence result (3.9), the second term on the right-hand side converges to zero in $L^1(\mu^{u,x})$. By the Cauchy-Schwarz inequality, Corollary 3.7, and the convergence result (3.10), the first term on the right-hand side converges to zero in $L^1(\mu^{u,x})$. This completes the proof. \square

COROLLARY 3.11. *The first variation $\Phi^{\sigma,x}(u; \varphi)$ satisfies*

$$(3.24) \quad \Phi^{\sigma,x}(u; \varphi) = E^{u,x} [K^{\sigma,u,1} M_\tau^\varphi].$$

In particular, the first variation $\Phi^{\sigma,x}(u; \varphi)$ vanishes if and only if $K^{\sigma,u,1}$ and M_τ^φ are uncorrelated with respect to $\mu^{u,x}$.

Proof. The first assertion follows from the definition (2.12) of $K^{\sigma,u,\alpha}$ and the fact that $E^{u,x}[\langle M^u, M^\varphi \rangle_\tau] = E^{u,x}[M_\tau^u M_\tau^\varphi]$. The second assertion follows from (3.24). \square

Corollary 3.11 is consistent with the assertion made earlier, that the optimal control that solves the stochastic optimal control problem (2.14) yields an optimal importance sampling measure, i.e. a zero-variance estimator for the value function F for any given pair of arguments.

3.4. First variation of first variation. In this section, we derive an expression for the first variation of the functional $\Phi^{\sigma,x}(u; \varphi)$ along an admissible perturbation ζ , where the argument u is perturbed and the parameters σ , x , and φ are held fixed. As in §3.3, we shall first define for a given sequence $(\delta_n)_n$ a sequence of finite differences:

$$(3.25) \quad h_n^{\sigma,u,\varphi,\zeta} := \left(K^{\sigma,u,0} M_\tau^\varphi + \frac{1}{\sigma} \langle M^u, M^\varphi \rangle_\tau \right) \frac{\mathcal{E}(M^{\delta_n \zeta})_\tau - 1}{\delta_n} + \frac{1}{\sigma} \left(\langle M^u, M^\zeta \rangle_\tau M_\tau^\varphi + \langle M^\varphi, M^\zeta \rangle_\tau + \frac{\delta_n}{2} \langle M^\zeta \rangle_\tau \right) \mathcal{E}(M^{\delta_n \zeta})_\tau.$$

The following result is analogous to Lemma 3.9:

LEMMA 3.12. *Suppose that Assumption 3.1 holds, that $u \in C^\alpha(\overline{\Omega})$, that $\varphi, \zeta \in L^\infty(\Omega)$, and that $(\delta_n)_n$ is an arbitrary nonzero sequence satisfying $\delta_n \rightarrow 0$ as $n \rightarrow \infty$. Let $(h_n^{\sigma,u,\varphi,\zeta})_n$ be the corresponding sequence given by (3.25). Then $(h_n^{\sigma,u,\varphi,\zeta})_n \subset L^1(\mu^{u,x})$ and*

$$\frac{\Phi^{\sigma,x}(u + \delta_n \zeta; \varphi) - \Phi^{\sigma,x}(u; \varphi)}{\delta_n} = E^{u,x} [h_n^{\sigma,u,\varphi,\zeta}].$$

Proof. We first prove that the terms on the first row of (3.25) belong to $L^1(\mu^{u,x})$. By the triangle inequality and the inequality (3.16) of arithmetic and geometric means, it holds that

$$\begin{aligned} & \left(K^{\sigma,u,0} M_\tau^\varphi + \frac{1}{\sigma} \langle M^u, M^\varphi \rangle_\tau \right) \frac{\mathcal{E}(M^{\delta_n \zeta})_\tau - 1}{\delta_n} \\ & \leq |K^{\sigma,u,0} M_\tau^\varphi|^2 + 2 \left| \frac{\mathcal{E}(M^{\delta_n \zeta})_\tau - 1}{\delta_n} \right|^2 + \sigma^{-2} \langle M^u, M^\varphi \rangle_\tau^2. \end{aligned}$$

The third term on the right-hand side belongs to $L^1(\mu^{u,x})$ by (3.18), the second term on the right-hand side belongs to $L^1(\mu^{u,x})$ by (3.14), and the first term on the right-hand side belongs to $L^1(\mu^{u,x})$, by statement (ii) of Corollary 3.5 and 3.7.

Of the terms on the second row of (3.25), we only need to show that the product $\langle M^u, M^\zeta \rangle_\tau M_\tau^\varphi \in L^2(\mu^{u,x})$, since the other terms were shown to belong to $L^1(\mu^{u,x})$ in the proof of Lemma 3.9. By the Cauchy-Schwarz inequality, (3.16), and (3.17),

$$(3.26) \quad E^{u,x} \left[(\langle M^u, M^\zeta \rangle_\tau M_\tau^\varphi)^2 \right] \leq (E^{u,x} [\langle M^u \rangle_\tau^4 + \langle M^\zeta \rangle_\tau^4] E^{u,x} [\langle M^\varphi \rangle_\tau])^{1/2},$$

which completes the proof that $h_n^{\sigma,u,\varphi,\zeta} \in L^1(\mu^{u,x})$ for all n .

The second conclusion of the lemma follows in the same manner as for the second conclusion of Lemma 3.9, i.e. by expanding the quadratic variation according to (3.19), using the reweighting formula (2.6), and taking differences. \square

We define the limiting random variable of the sequence $(h_n^{\sigma,u,\varphi,\zeta})_n$ by

$$(3.27) \quad h^{\sigma,u,\varphi,\zeta} := K^{\sigma,u,0} M_\tau^\varphi M_\tau^\zeta + \frac{1}{\sigma} (\langle M^\varphi, M^\zeta \rangle_\tau + M_\tau^\zeta \langle M^u, M^\varphi \rangle_\tau + M_\tau^\varphi \langle M^u, M^\zeta \rangle_\tau),$$

and show that, in expectation, $h^{\sigma,u,\varphi,\zeta}$ equals the first variation of $\Phi^{\sigma,x}(u; \varphi)$ along ζ :

THEOREM 3.13. *If the hypotheses of Lemma 3.12 are satisfied, then $h^{\sigma,u,\varphi,\zeta}$ belongs to $L^1(\mu^{u,x})$, $h_n^{\sigma,u,\varphi,\zeta}$ converges to $h^{\sigma,u,\varphi,\zeta}$ in $L^1(\mu^{u,x})$, and the first variation of $\Phi^{\sigma,x}(u; \varphi)$ along ζ satisfies*

$$(3.28) \quad \Phi^{\sigma,x}(u; \varphi, \zeta) := \lim_{n \rightarrow \infty} \frac{\Phi^{\sigma,x}(u + \delta_n \zeta; \varphi) - \Phi^{\sigma,x}(u; \varphi)}{\delta_n} = E^{u,x} [h^{\sigma,u,\varphi,\zeta}].$$

Proof. As in the proof of Theorem 3.10, the third conclusion follows from the second, by (3.22). The first conclusion, that $h^{\sigma,u,\varphi,\zeta}$ belongs to $L^1(\mu^{u,x})$ follows from Corollary 3.7, (3.18), and (3.26).

It remains to prove the second conclusion, i.e. the convergence of $h_n^{\sigma,u,\varphi,\zeta}$ to $h^{\sigma,u,\varphi,\zeta}$ in $L^1(\mu^{u,x})$. We apply the triangle inequality to obtain

$$\begin{aligned} |h_n^{\sigma,u,\varphi,\zeta} - h^{\sigma,u,\varphi,\zeta}| &\leq (|K^{\sigma,u,0} M_\tau^\varphi| + \sigma^{-1} |\langle M^u, M^\varphi \rangle_\tau|) \left| \frac{\mathcal{E}(M^{\delta_n \zeta})_\tau - 1}{\delta_n} - M_\tau^\zeta \right| \\ &\quad + \frac{1}{\sigma} (|\langle M^\varphi, M^\zeta \rangle_\tau| + |\langle M^u, M^\zeta \rangle_\tau M_\tau^\varphi|) |\mathcal{E}(M^{\delta_n \zeta})_\tau - 1| \\ &\quad + \frac{|\delta_n|}{2\sigma} \langle M^\zeta \rangle_\tau \mathcal{E}(M^{\delta_n \zeta})_\tau. \end{aligned}$$

The $L^1(\mu^{u,x})$ convergence to zero of the term on the third row of the right-hand side was shown in the proof of Theorem 3.10. Of the two terms on the second row, only the convergence of the product of $|\langle M^u, M^\zeta \rangle_\tau M_\tau^\varphi|$ and $|\mathcal{E}(M^{\delta_n \zeta})_\tau - 1|$ has not been proven, but this follows by the Cauchy-Schwarz inequality, (3.26), and (3.9). The $L^1(\mu^{u,x})$ convergence to zero of the terms on the first row follows from the Cauchy-Schwarz inequality and the convergence result (3.10). \square

We now use Theorem 3.13, along with some standard results from the calculus of variations concerning second-order conditions for convexity of C^2 functionals defined on Banach spaces, in order to specify a sufficient condition for convexity of the control functional ϕ defined in (2.13).

THEOREM 3.14. *Assume that the hypotheses of Theorem 3.13 hold. Let $\sigma > 0$ and $x \in \Omega \setminus \partial\Omega$ be fixed. If the terminal cost κ_t satisfies the lower bound*

$$(3.29) \quad \kappa_t(y) \geq \sigma^{-1} \quad \forall y \in \partial\Omega,$$

then the undiscretised control functional $\phi^{\sigma,x} : C^{1,\alpha}(\overline{\Omega}; \mathbb{R}^d) \rightarrow \mathbb{R}$ is strictly convex.

Proof. By setting $\varphi = \zeta$ in (3.27) and taking expectations, we obtain the following expression for the second variation of $\phi^{\sigma,x}$ at u in the direction of φ :

$$(3.30) \quad \Phi^{\sigma,x}(u; \varphi, \varphi) = E^{u,x} [K^{\sigma,u,0}(M_\tau^\varphi)^2 + \sigma^{-1} (\langle M^\varphi \rangle_\tau + 2\langle M^u, M^\varphi \rangle_\tau M_\tau^\varphi)].$$

Recall (see, e.g. [28, Section 2.9, Corollary 4]) that if the second variation of a C^2 functional on a Banach space is positive definite, then the functional is strictly convex. With this in mind, we show that (3.29) suffices to guarantee positive definiteness of the second variation. Let $u \in C^{1,\alpha}(\bar{\Omega})$ and $0 \neq \varphi \in C^{1,\alpha}(\bar{\Omega})$ be arbitrary. Observe that, by the Kunita-Watanabe inequality and Young's inequality, we have

$$|\langle M^u, M^\varphi \rangle_\tau M_\tau^\varphi| \leq (\langle M^u \rangle_\tau \langle M^\varphi \rangle_\tau)^{1/2} |M_\tau^\varphi| \leq 4^{-1} \langle M^u \rangle_\tau (M_\tau^\varphi)^2 + \langle M^\varphi \rangle_\tau.$$

Therefore, we have

$$(3.31) \quad \begin{aligned} \Phi^{\sigma,x}(u; \varphi, \varphi) &\geq E^{u,x} \left[K^{\sigma,u,0}(M_\tau^\varphi)^2 - \frac{2}{\sigma} \left(\frac{1}{4} \langle M^u \rangle_\tau (M_\tau^\varphi)^2 + \langle M^\varphi \rangle_\tau \right) \right] \\ &= E^{u,x} [(W - \sigma^{-1})(M_\tau^\varphi)^2], \end{aligned}$$

where we used the definition (2.12) of $K^{\sigma,u,\alpha}$ and the Itô isometry to obtain (3.31). Since $x \in \Omega \setminus \partial\Omega$ and κ_r is nonnegative by Assumption 3.1, it follows from (3.29) that $W - \sigma^{-1}$ is strictly positive $\mu^{u,x}$ -almost surely. Furthermore, since $\varphi \neq 0$ and since $x \in \Omega \setminus \partial\Omega$, it follows that $(M_\tau^\varphi)^2 = \langle M^\varphi \rangle_\tau$ is strictly positive $\mu^{u,x}$ -almost surely. This proves that the random variable inside the expectation in (3.31) is strictly positive $\mu^{u,x}$ -almost surely, and thus the variation $\Phi^{\sigma,x}(u; \varphi, \varphi)$ of the control functional $\phi^{\sigma,x}$ is positive definite on $C^{1,\alpha}(\bar{\Omega})$. \square

Remark on second-order characterisation of convexity. The advantage of proving strict convexity of the control functional $\phi^{\sigma,x}$ by proving positive definiteness of the second variation is that one can bound the second variation from below by a quantity that does not contain any exponential martingales, as we have shown above. This leads to a simple sufficient condition (3.29) for positive definiteness. On the other hand, one can verify using the reweighting formula that the zeroth- and first-order characterisations of convexity involve expressions that contain exponential martingale terms, and thus the sufficient conditions obtained from these characterisations are more complicated. Furthermore, for the zeroth-order characterisation, one needs to check that the Radon-Nikodym derivative exists and is a uniformly integrable exponential martingale, in order to express all the expected values that appear in the zeroth-order characterisation in terms of a common path measure $\mu^{u',x}$. Although Novikov's condition is well-known in stochastic analysis for being a sufficient condition for uniform integrability, we do not know of any sufficient conditions for Novikov's condition that can be verified for arbitrary bounded domains of sufficient regularity and stopping times that are integrable but not almost-surely bounded, aside from the hypotheses of Theorem 3.4. However, once the hypotheses of Theorem 3.4 hold, then it follows that the second variation exists, and one can prove the strict convexity of the control functional using the second-order characterisation, as we have done above.

Given that we have imposed strong conditions on the data of the stochastic optimal control problem, such as boundedness and regularity of the domain Ω and restriction to feedback controls belonging to $C^{1,\alpha}(\bar{\Omega}; \mathbb{R}^d)$, it seems reasonable that we obtain strong results on the control functional. Let $\Pi_2(B)$ be the set of all predictable processes that are adapted to the filtration generated by the Brownian motion B . Given

$u \in C^\alpha(\overline{\Omega}; \mathbb{R}^d)$, define the function $\|\cdot\|_{B,u} : \Pi_2(B) \rightarrow [0, \infty)$ according to

$$(3.32) \quad \|\varphi\|_{B,u} := (E^{u,x}[\langle M^\varphi \rangle_\tau])^{1/2}.$$

It follows from the definition of the quadratic variation process that $\|\cdot\|_{B,u}$ defines a norm on $\Pi_2(B)$. Since feedback processes $u_t := u(X_t^u)$ are predictable processes, it follows that $\|\cdot\|_{B,u}$ induces a norm on the set of feedback control processes, including those generated by feedback controls belonging to $C^{1,\alpha}(\overline{\Omega}; \mathbb{R}^d)$. We can therefore consider $\|\cdot\|_{B,u}$ as a norm on $C^{1,\alpha}(\overline{\Omega}; \mathbb{R}^d)$. We then have the following

COROLLARY 3.15. *Suppose that*

$$(3.33) \quad \gamma := \gamma(\sigma, \kappa_t) := \min_{y \in \partial\Omega} \kappa_t(y) - \sigma^{-1}$$

is strictly positive. Then for all $u \in C^{1,\alpha}(\overline{\Omega}; \mathbb{R}^d)$, the second variation $\Phi^{\sigma,x}(u; \varphi, \varphi)$ is a coercive functional of φ with parameter γ on $C^{1,\alpha}(\overline{\Omega}; \mathbb{R}^d)$, in the sense that

$$(3.34) \quad \Phi^{\sigma,x}(u; \varphi, \varphi) \geq \gamma(\sigma, \kappa_t) \|\varphi\|_{B,u}^2$$

for all $0 \neq \varphi \in C^{1,\alpha}(\overline{\Omega}; \mathbb{R}^d)$.

We emphasise that Theorem 3.14 are consistent with the classical result (Theorem 2.2) concerning the existence of a unique solution to the linear Dirichlet problem. In the context of linear elliptic Dirichlet boundary value problems, κ_t is a continuous function on a compact set. Therefore, one can always ensure that the sufficient condition (3.29) holds, by considering the translated terminal cost function $\hat{\kappa}_t := \kappa_t + C$, where $C = \max\{0, -\gamma(\sigma, \kappa_t)\}$.

4. Approximation by a finite-dimensional subspace. In this section, we consider some consequences of the strict convexity result Theorem 3.14, given that Assumption 3.1 holds. The first consequence of Theorem 3.14 that we shall consider (Theorem 4.4 below) concerns the well-posedness of the gradient descent algorithm proposed by Hartmann and Schütte in [13]. The second consequence concerns the best approximation of the value function $F(\sigma, \cdot)$ that solves the Hamilton-Jacobi-Bellman equation of the stochastic optimal control problem (2.14).

Fix a finite collection $b := (b_i)_{i=1}^n$ of linearly independent, nonconstant functions $b_i \in C^{2,\alpha}(\overline{\Omega}; \mathbb{R})$, and let $S_i := \{x \in \overline{\Omega} \mid b_i(x) \neq 0\}$ denote the support of b_i in $\overline{\Omega}$. We shall assume that

$$(4.1) \quad \lambda(S_i) > 0, \quad \cup_{i=1}^n \overline{S_i} = \overline{\Omega},$$

i.e. the support of each function b_i has strictly positive Lebesgue measure, and the closures of the supports cover $\overline{\Omega}$. Let $\varphi_i := \nabla b_i$ and consider the following subspace of $C^{1,\alpha}(\overline{\Omega}; \mathbb{R}^d)$:

$$(4.2) \quad \mathcal{U}_{\text{Markov}}(b) := \left\{ u \in \mathcal{U}_{\text{Markov}} \mid u = u^a := \sum_i a_i \nabla b_i \text{ for some } a \in \mathbb{R}^n \right\}.$$

By linear independence of the basis vector fields $(\varphi_i)_i$, it holds that every $u \in \mathcal{U}_{\text{Markov}}(b)$ admits a unique representation $u = u^a$. Therefore, the restriction of the control functional $\phi^{\sigma,x}$ to the subspace $\mathcal{U}_{\text{Markov}}(b)$ may be uniquely represented by

$$(4.3) \quad \hat{\phi}^{\sigma,x} : \mathbb{R}^n \rightarrow \mathbb{R}, \quad \hat{\phi}^{\sigma,x}(a) := \phi^{\sigma,x}(u^a).$$

By Corollary 3.11, we obtain an expression for the partial derivative of the function $\hat{\phi}^{\sigma,x}$ with respect to the i -th coordinate a_i .

4.1. An initial value problem. Consider the initial value problem

$$(4.4) \quad \frac{d}{dt}a(s) = -\nabla_a \phi^{\sigma,x}(a(s)), \quad a(0) \in \mathbb{R}^n.$$

The main result of this section concerns the existence of solutions to (4.4). To prove the result, we make the following

ASSUMPTION 4.1. *Let $\nabla V \in C^\alpha(\overline{\Omega})$ and τ be defined as in (2.7). Then for all $1 \leq i \leq n$, there exists a pair $0 \leq t_0(i) < t_1(i) \leq \tau$ such that for all $t_0(i) < t < t_1(i)$, $X_t^0 \in S_i$ with positive P^x -probability.*

PROPOSITION 4.2. *Suppose that Assumptions 3.1 and 4.1 hold, and let τ be defined as in (2.7). If $u^a \in \mathcal{U}_{\text{Markov}}(b)$, then for all $1 \leq i \leq n$, there exists a pair $0 \leq t_0(i) < t_1(i) \leq \tau$ such that for all $t_0(i) < t < t_1(i)$, $X_t^{u^a} \in S_i$ with positive P^x -probability.*

Proof. By Assumption 3.1, we may apply Corollary 3.5 to assert that $\mathcal{E}((M^{u^a})^\tau)$ is a $L^2(\mu^{0,x})$ -bounded martingale, and hence is uniformly integrable with respect to $\mu^{0,x}$. By the change of measure theorem, the restrictions of $\mu^{0,x}$ and $\mu^{u^a,x}$ to \mathcal{F}_t are mutually absolutely continuous for all $t > 0$. The desired conclusion follows from Assumption 4.1. \square

COROLLARY 4.3. *Given an arbitrary $a(0) \in \mathbb{R}^n$ and compact neighbourhood \overline{N} of $a(0)$, it holds that*

$$(4.5) \quad m^x(\overline{N}) := \min_{a \in \overline{N}} \min_{1 \leq i \leq n} E^{u^a,x}[\langle M^{\varphi_i} \rangle_\tau] > 0.$$

Proof. By the representation (4.2), the distribution $\mu^{u^a,x}$ depends continuously on a . Therefore the map $a \mapsto \min_i E^{u^a,x}[\langle M^{\varphi_i} \rangle_\tau]$ is continuous. By Proposition 4.2, the map is strictly positive for every $a \in \mathbb{R}^n$. The conclusion follows by compactness. \square

Now consider the restriction to $\mathcal{U}_{\text{Markov}}(b)$ of the stochastic optimal control problem defined in (2.14),

$$(4.6) \quad \min_{a \in \mathbb{R}^n} \widehat{\phi}^{\sigma,x}(a) \quad \text{subject to (2.1)}.$$

The restricted problem differs from the original stochastic optimal control problem (2.14) only in the function space over which the control functional is to be minimised. The gradient descent algorithm proposed in [13] for solving the restricted problem above consisted of the following steps: Given $\mathcal{U}_{\text{Markov}}(b)$, an initial condition $a(0)$, $N_{\text{iter}}, N_{\text{traj}} \in \mathbb{N}$, and a sequence of descent steps $(h_\ell)_{0 \leq \ell \leq N_{\text{iter}}-1} \subset (0, \infty)$, we have:

- for** $\ell = 1 : N_{\text{iter}}$ **do**
- Sample N_{traj} trajectories of (2.1) with $u = u^{a(i-1)}$
 - Approximate $\widehat{\phi}^{\sigma,x}(a(\ell-1))$ and $\nabla_a \widehat{\phi}^{\sigma,x}(a(\ell-1))$ by Monte Carlo
 - Update the coefficient vector $a(\ell-1)$ to $a(\ell)$ via

$$(4.7) \quad a(\ell) = a(\ell-1) - h_{\ell-1} \nabla_a \widehat{\phi}^{\sigma,x}(a(\ell-1))$$

end

That is, the gradient descent algorithm (4.7) is the discrete time, finite-sample size analogue of the initial value problem (4.4). We shall study the well-posedness of

the gradient descent algorithm via (4.4). The following result shows that, under Assumptions 3.1 and 4.1, the restricted problem (4.6) is well-posed.

THEOREM 4.4. *Suppose that the hypotheses of Theorem 3.13 are satisfied. If (3.29) holds, then there exists a unique $a_\infty \in \mathbb{R}^n$ such that*

$$(4.8) \quad 0 = \nabla_a \hat{\phi}^{\sigma, x}(a_\infty),$$

and for every initial condition $a(0) \in \mathbb{R}^n$, there exists a unique solution $(a(t))_{0 \leq t < \infty}$ of the initial value problem (4.4) to a_∞ . If, in addition to the above hypotheses, it holds that

$$(4.9) \quad \overline{S_i} \cap \overline{S_j} = \partial S_i \cap \partial S_j,$$

then the unique solution converges exponentially to a_∞ , with

$$(4.10) \quad |a_\infty - a(t)|^2 \leq |a(0) - a_\infty|^2 \exp(-t\gamma m^x(\overline{N})),$$

where $m^x(\overline{N})$ is defined in (4.5).

Proof. Since the hypotheses of Theorem 3.13 hold, and since (3.29) holds, it follows that the control functional $\phi^{\sigma, x}$ is strictly convex on $C^{1, \alpha}(\overline{\Omega}; \mathbb{R}^d)$, by Theorem 3.14. This implies that the function $\hat{\phi}^{\sigma, x}$ is twice continuously differentiable and strictly convex on any bounded, open, nonempty convex subset of \mathbb{R}^n , so there exists a unique global minimiser that satisfies (4.8). By the standard results on existence and uniqueness of solutions to ordinary differential equations, it follows that for any initial condition $a(0)$, there exists a unique solution $(a(t))_{t \geq 0}$ that exists for all time.

By Theorem 3.13, (3.27), and linearity of the representation u^a , it holds that, given the Hessian $\nabla_a^2 \hat{\phi}^{\sigma, x} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times n}$, arbitrary $a' \in \mathbb{R}^n$ and $0 \neq z \in \mathbb{R}^n$,

$$z \cdot \nabla_a^2 \hat{\phi}^{\sigma, x}(a')z = E^{u^{a'}, x} \left[K^{\sigma, u^{a'}, 0}(M_\tau^{u^z})^2 + \sigma^{-1} \left(\langle M^{u^z} \rangle_\tau + 2 \langle M^{u^{a'}} \rangle_\tau M_\tau^{u^z} \right) \right].$$

By (3.30), the value of $z \cdot \nabla_a^2 \hat{\phi}^{\sigma, x}(a')z$ equals the value of the second variation of $\phi^{\sigma, x}$ at $u^{a'}$ along the element u^z . By Corollary 3.15,

$$(4.11) \quad z \cdot \nabla_a^2 \hat{\phi}^{\sigma, x}(a')z \geq \gamma E^{u^{a'}, x} \left[\langle M^{u^z} \rangle_\tau \right].$$

By the chain rule, Taylor's theorem and the property (4.8), there is a $z \in \mathbb{R}^n$ that is a convex combination of $a(s)$ and a_∞ such that

$$(4.12) \quad \begin{aligned} \frac{d}{dt} \frac{1}{2} |a(s) - a_\infty|^2 &\leq -\frac{1}{2} (a_\infty - a(s)) \cdot \nabla_a^2 \hat{\phi}^{\sigma, x}(z) (a_\infty - a(s)) \\ &\leq -\frac{1}{2} \gamma E^{u^z, x} \left[\langle M^{u^{a_\infty - a(s)}} \rangle_\tau \right]. \end{aligned}$$

Observe that by Assumption 3.1, $\mu^{u^a, x}$ is unique, and therefore X^{u^a} has the strong Markov property. Under the condition (4.9) on the supports, it holds that $\varphi_i(x) \cdot \varphi_j(x)$ is zero except on a set of zero Lebesgue measure, and therefore $\langle M^{\varphi_i}, M^{\varphi_j} \rangle_\tau$ has vanishing $\mu^{u^a, x}$ -expectation. Let \overline{N} be the smallest compact set that contains all convex combinations of $a(t)$ and a_∞ for all $t \geq 0$. It follows from the vanishing of the covariance processes that $\phi^{\sigma, x}$ is strongly convex on N , since

$$E^{u^z, x} \left[\langle M^{u^{a_\infty - a(s)}} \rangle_\tau \right] \geq m^x(\overline{N}) |a_\infty - a(s)|^2.$$

Substituting the above inequality in (4.12), integrating the resulting differential inequality, and applying Gronwall's lemma, we obtain (4.10). \square

4.2. Approximation of solution to the HJB equation. In this section, we relate the unique minimiser of the restricted control functional $\widehat{\phi}^{\sigma,x}$ to the best approximation in the finite-dimensional subspace $\mathcal{U}_{\text{Markov}}(b)$ of $F(\sigma, \cdot)$. Recall that, by Theorem 2.2, the function $\psi(\sigma, \cdot)$ for ψ defined in (2.9a) belongs to the space $C^{2,\alpha}(\overline{\Omega})$. Since $\psi(\sigma, \cdot) = \exp(-\sigma F(\sigma, \cdot))$, this implies that $F(\sigma, \cdot) \in C^{2,\alpha}(\overline{\Omega})$, and since $\overline{\Omega}$ is compact, this implies that $F(\sigma, \cdot) \in L^2(\Omega; \mathbb{R})$. In fact, by the same reasoning, $F(\sigma, \cdot)$ is an element of the Sobolev space $H^{2,p}(\Omega)$ for $1 \leq p < \infty$. The unique equilibrium a_∞ defined in (4.8) gives the best approximation in the space $\mathcal{U}_{\text{Markov}}(b)$ of the optimal control u_{opt}^σ . Therefore, the unique equilibrium a_∞ gives the best approximation of $\nabla_x F(\sigma, \cdot)$, up to scaling. Since we have assumed Ω to be a connected, bounded domain with $C^{2,\alpha}$ -boundary, we may apply the Poincaré inequality as follows: define

$$F_{\text{approx}}(\sigma, x') := \sum_i (a_\infty)_i b_i(x') + \left(\widehat{\phi}^{\sigma,x}(a_\infty) - \sum_i (a_\infty)_i b_i(x) \right),$$

where the additive constant in the parentheses on the right-hand side ensures that $F_{\text{approx}}(\sigma, x) = \widehat{\phi}^{\sigma,x}(a_\infty)$. Define F_{err} and the average value $(F_{\text{err}})_\Omega$ by

$$F_{\text{err}}(\sigma, x) := F(\sigma, x) - F_{\text{approx}}(\sigma, x), \quad (F_{\text{err}})_\Omega := \frac{1}{\text{vol}(\Omega)} \int_\Omega F_{\text{err}}(\sigma, x) dx.$$

Then since $F_{\text{err}}(\sigma, \cdot) \in W^{1,p}(\Omega)$, the Poincaré inequality yields the estimate

$$(4.13) \quad \|F_{\text{err}}(\sigma, \cdot) - (F_{\text{err}})_\Omega\|_{L^p(\Omega)} \leq C(n, p, \Omega) \|\nabla_x F_{\text{err}}(\sigma, \cdot)\|_{L^p(\Omega)}$$

for $1 \leq p < \infty$ and $C(n, p, \Omega)$ independent of $F_{\text{err}}(\sigma, \cdot)$. The estimate emphasises the importance of the approximation quality of the subspace $\mathcal{U}_{\text{Markov}}(b)$: as the approximation quality of $\mathcal{U}_{\text{Markov}}(b)$ improves, the right-hand side of (4.13) decreases, giving smaller L^p -errors in the approximation of $F(\sigma, \cdot)$.

5. Conclusion. In this article, we proved that the control functional of a stochastic optimal control problem is strictly convex, under the conditions that guarantee the existence and uniqueness of a value function that is uniformly Hölder continuous on the closure of a bounded, $C^{2,\alpha}$ domain. Our results relied on applying Fredholm theory in order to show that the moment generating function of the first exit time with respect to the law $\mu^{u,x}$ for feedback controls $u \in L^\infty(\Omega)$ is finite on any bounded open interval containing the origin. From this result, we obtained the L^2 -boundedness of the exponential martingale associated to the feedback control. We then used L^2 -boundedness to derive an expression for the second variation of the control functional. Positive definiteness then followed from a specific lower bound on the terminal cost function. The result of strict convexity holds independently of the dimension d of the domain. The significance of this result is that one may solve the Hamilton-Jacobi-Bellman boundary value problem, and thereby the rare event importance sampling problem, by using methods for convex optimisation in conjunction with Monte Carlo methods. The advantage of doing so is that both classes of methods scale well with dimension.

Acknowledgments. The author thanks Christof Schütte, Carsten Hartmann, and Tim Sullivan from the Freie Universität Berlin for their encouragement and support of the thesis on which this article is based, Patrick J. Fitzsimmons from the University of California at San Diego for pointing out Kazamaki's book [17], and Giacomo di Gesù from CERMICS - École des Ponts ParisTech for pointing out an important subtlety concerning the mutual absolute continuity of path measures.

REFERENCES

- [1] R. J. ALLEN, D. FRENKEL, AND P. R. TEN WOLDE, *Simulating rare events in equilibrium or nonequilibrium stochastic systems*, J. Chem. Phys., 124 (2006).
- [2] S. ASMUSSEN, D. P. KROESE, AND R. Y. RUBINSTEIN, *Heavy tails, importance sampling and cross-entropy*, Stochastic Models, 21 (2005), pp. 57–76.
- [3] V. S. BORKAR, *A convex analytic approach to markov decision processes*, 78 (1988), pp. 583–602.
- [4] P. DAI PRA, L. MENEGHINI, AND W. J. Runggaldier, *Connections between stochastic control and dynamic games*, Math. Control Signals Syst., 9 (1996), pp. 303–326.
- [5] E. DARVE, D. RODRIGUEZ-GOMEZ, AND A. POHORILLE, *Adaptive biasing force method for scalar and vector free energy calculations*, J. Chem. Phys., 128 (2008).
- [6] C. DELLAGO, P. G. BOLHUIS, F. S. CSAJKA, AND D. CHANDLER, *Transition path sampling and the calculation of rate constants*, J. Chem. Phys., 108 (1998), pp. 1964–1977.
- [7] P. DUPUIS AND W. M. MCENEANEY, *Risk-sensitive and robust escape criteria*, SIAM J. Control Optim., 35 (1997), pp. 2021–2049.
- [8] P. DUPUIS, K. SPILIOPOULOS, AND H. WANG, *Importance sampling for multiscale diffusions*, Multiscale Model. Simul., 10 (2012), pp. 1–27.
- [9] P. DUPUIS AND H. WANG, *Importance sampling, large deviations, and differential games*, Stochastics Stochastics Rep., 76 (2004), pp. 481–508.
- [10] ———, *Dynamic importance sampling for uniformly recurrent markov chains*, Ann. Appl. Probab., 15 (2005), pp. 1–38.
- [11] D. GILBARG AND N. S. TRUDINGER, *Elliptic Partial Differential Equations of Second Order*, Grundlehren der mathematischen Wissenschaften, Springer, Berlin, 2001. (Reprint of the 1998 edition).
- [12] C. HARTMANN, R. BANISCH, M. SARICH, T. BADOWSKI, AND C. SCHÜTTE, *Characterization of rare events in molecular dynamics*, Entropy, 16 (2013), pp. 350–376.
- [13] C. HARTMANN AND C. SCHÜTTE, *Efficient rare event simulation by optimal nonequilibrium forcing*, J. Stat. Mech. Theor. Exp., 2012 (2012), p. P11004.
- [14] C. JARZYNSKI, *Nonequilibrium equality for free energy differences*, Phys. Rev. Lett., 78 (1997), pp. 2690–2693.
- [15] ———, *Rare events and the convergence of exponentially averaged work values*, Phys. Rev. E, 73 (2006), p. 046105.
- [16] I. KARATZAS AND S. E. SHREVE, *Brownian Motion and Stochastic Calculus*, Springer, Berlin, 1991.
- [17] N. KAZAMAKI, *Continuous Exponential Martingales and BMO*, Lecture Notes in Mathematics, Springer-Verlag, Berlin, 1994.
- [18] S. KULLBACK, *Information Theory and Statistics*, Wiley, New York, 1959.
- [19] A. LAIO AND M. PARRINELLO, *Escaping free-energy minima*, Proc. Natl. Acad. Sci. U.S.A., 99 (2002), pp. 12562–12566.
- [20] T. LELIÈVRE, M. ROUSSET, AND G. STOLTZ, *Free Energy Computations: A Mathematical Perspective*, Imperial College Press, London, 2010.
- [21] E. J. MCSHANE, *Jensen's inequality*, Bull. Amer. Math. Soc., 43 (1937), pp. 521–527.
- [22] D. PASSERONE AND M. PARRINELLO, *Action-derived molecular dynamics in the study of rare events*, Phys. Rev. Lett., 87 (2001), p. 108302.
- [23] D. REVUZ AND M. YOR, *Continuous Martingales and Brownian Motion*, Grundlehren der mathematischen Wissenschaften, Springer, Berlin, 2005. Corrected third printing of third edition.
- [24] R. Y. RUBINSTEIN, *Optimization of computer simulation models with rare events*, Eur. J. Oper. Res., 99 (1997), pp. 89 – 112.
- [25] E. VANDEN-EIJNDEN AND J. WEARE, *Rare event simulation of small noise diffusions*, Comm. Pure Appl. Math., 65 (2012), pp. 1770–1803.
- [26] A. F. VOTER, *Hyperdynamics: Accelerated molecular dynamics of infrequent events*, Phys. Rev. Lett., 78 (1997), pp. 3908–3911.
- [27] ———, *A method for accelerating the molecular dynamics simulation of infrequent events*, J. Chem. Phys., 106 (1997), pp. 4665–4677.
- [28] E. ZEIDLER, *Applied Functional Analysis: Main Principles and Their Applications*, vol. 109 of Applied Mathematical Sciences, Springer-Verlag, Berlin, 1995.
- [29] W. ZHANG, H. WANG, C. HARTMANN, M. WEBER, AND C. SCHÜTTE, *Applications of the cross-entropy method to importance sampling and optimal control of diffusions*, SIAM J. Sci. Comput., 36 (2014), pp. A2654–A2672.